PCT/GB 2004 / 0 0 4 2 2 2

INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

PRIORITY DOCUMENT

SUBMITTED OR TRANSMITTED IN COMPLIANCE WITH RULE 17.1(a) OR (b)

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

REC'D 30 NOV 2004
WIPO PCT

Signed

Dated 17 November 2004

BEST AVAILABLE COPY

Patents Form 1/77

Patents Act 1977
(Rule 16)

The Patent Office

1/77

07OCT03 E842691-1 D00358
P01/7700 0.00-0323429.1

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

THE PATENT OFFICE
RM
- 7 OCT 2003
RECEIVED BY FAX

The Patent Office

Cardiff Road
Newport
South Wales
NP10 8QQ

1.	Your reference	2003P15450 GB / R76 / JJP / RRI
2.	Patent application number *(The Patent Office will fill in this part)*	0323429.1 - 7 OCT 2003
3.	Full name, address and postcode of the or of each applicant *(underline all surnames)*	ROKE MANOR RESEARCH LIMITED Roke Manor Old Salisbury Lane, Romsey Hampshire SO51 0ZN
	Patents ADP number *(if you know it)*	5615455007
	If the applicant is a corporate body, give the country/state of its incorporation	UNITED KINGDOM
4.	Title of the invention	Spectrum Sharing
5.	Name of your agent *(if you have one)*	Jan Payne
	"Address for service" in the United Kingdom to which all correspondence should be sent *(including the postcode)*	Siemens plc Intellectual Property Department The Lodge, Roke Manor Romsey, Hampshire SO51 0ZN
	Patents ADP number *(if you know it)*	853012300
6.	Priority: Complete this section if you are declaring priority from one or more earlier patent applications, filed in the last 12 months.	Country / Priority application number *(if you know it)* / Date of filing *(day / month / year)*
7.	Divisionals, etc: Complete this section only if this application is a divisional application or resulted from an entitlement dispute *(see note f)*.	

Patents Form 1/77

Patents Act 1977
(Rule 16)

The Patent Office

1/77

8. Is a Patents Form 7/77 (Statement of inventorship and of right to grant of a patent) required in support of this request? Answer YES if: *a) any applicant named in part 3 is not an inventor, or b) there is an inventor who is not named as an applicant, or c) any named applicant is a corporate body. Otherwise, answer NO (see note d).*	Yes
9. Accompanying documents: A patent application must include a description of the invention. Not counting duplicates, please enter the number of pages of each item accompanying this form.	
Continuation sheets of this form	
Description	9
Claim*(s)*	0
Abstract	0
Drawing*(s)*	0 only Dm
10. If you are also filing any of the following, state how many against each item.	
Priority documents	
Translation of priority documents	
Statement of inventorship and right to grant a patent *(Patents Form 7/77)*	
Request for preliminary examination and search *(Patents Form 9/77)*	
Request for substantive examination *(Patents Form 10/77)*	
Any other documents *(please specify)*	

I/We request the grant of a patent on the basis of this application

Signature | Date

Spectrum Sharing

The fundamental problem of spectrum sharing is not the transmitters but the receivers. It is possible to sense the presence or absence of signal that is already operating on a frequency one wishes to share. However, this does not determine whether transmission on that frequency will cause interference or not. The presence of strong signal might indicate that a remote receiver will have no problem with any interference your terminal might generate. Equally a weak signal might indicate that a nearby receiver cannot cope with any interference your terminal might produce. This runs contrary to the standard philosophy of dynamic channel assignment in which the channels with minimum interference are preferred. The problem is that we do not know what to do unless we know where the receivers are.

In order to solve this problem it is proposed to have receivers that transmit and transmitters that receive. Clearly they cannot do so on the same frequency, but if a nearby frequency is assigned for this purpose then this becomes possible. The compromise is to use a frequency that is near enough to have good correlation of path loss, but far enough to allow diplexing.

The basic principle is that some or all of the receivers transmit a beacon at a suitable power level. Any terminal that can hear a beacon related to a given frequency must reduce its transmitter power according to the strength of the beacon. If the permissible power is too low to allow the required communications then the terminal cannot use that frequency. A terminal operating dynamic channel allocation (DCA) would scan around the beacons and select the frequency whose beacon was received at the lowest power. If multiple beacons representing the same frequency are received at different powers the power of interest is taken to be the strongest.

We can arrange the beacon transmission band for a band of frequencies to be at either end (or both ends) of that band. We have one separate beacon signal relating to each frequency in the band and these must be multiplexed together.

The concept will expanded in the following sections...

1 Beacon Power and Interference Limitation

The situation is illustrated in Figure 1.

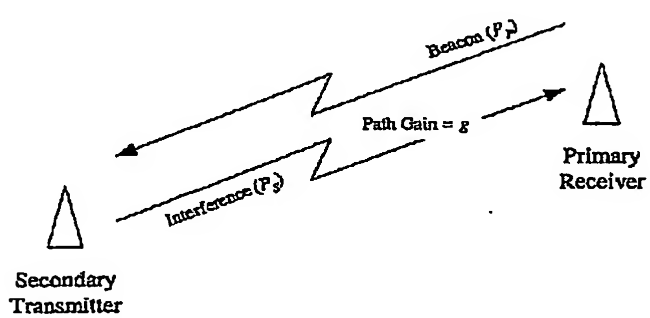

Figure 1 Beacon Transmission

We have two links, one with primary and one with secondary status. The transmitter for the primary link and the receiver for the secondary link are not shown. The primary receiver is transmitting a beacon in order to guarantee reception. We assume the following...

P_P is the primary receiver beacon power
P_S is the secondary transmitter power spectral density
g is the path gain between the two equipment connectors – this includes propagation path gain (<<1), antenna gains and feeder losses
N_P is the primary receiver noise figure
N_S is the noise figure of the beacon receiver at the secondary transmitter
B_S is the effective bandwidth at the beacon receiver for receiving the beacon
γ_S is the minimum signal to noise ratio for receiving the beacon in its effective bandwidth
d is the factor by which we allow the primary receiver to be de-sensitised

We need to set the level of P_S such that it is measurable at a point where a secondary transmitter could produce non-trivial interference at the primary receiver. We do this by fixing the maximum power that a secondary transmitter can emit when the beacon is undetectable such that only acceptable interference arises.

The limit sensitivity for the beacon receiver is $L_S = kTB_SN_S\gamma_S$ (where k is Boltzmann's constant and T is the operating temperature). Thus, if the received beacon power is less than or equal to L_S, then the secondary transmitter will emit a power spectral density of $P_{S\text{-}MAX}$ Watts/Hz. This is set so that if the beacon signal is equal to L_S, then only acceptable interference will be generated at the primary receiver. If we assume that this condition happens for a path gain of g then the received beacon power will be given by $L_S = P_P.g$. Thus $P_P.g = kTB_SN_S\gamma_S$.

The interference power spectral density generated at the primary receiver will then be $P_{S\text{-}MAX}.g$. The acceptable level for this will be $kTN_P(d\text{-}1)$. Thus $P_{S\text{-}MAX}.g = kTN_P(d\text{-}1)$. Eliminating g from the two equations gives

$$P_{S\text{-}MAX} = P_P(d-1)\frac{N_P}{\gamma_S B_S N_S}$$

This assumes that the operating temperature is the same at both locations which will be true to a reasonable approximation.

In practice, for universal operation, a reference would need to be set. We could, for example set the reference at 1W for a de-sensitisation of 3 dB with a 0 dB noise figure primary receiver. This would (conveniently) give a beacon power for 3 dB de-sensitisation equal to $1/N_P$ Watts.

As an example, consider the following case. Suppose...

- We allow de-sensitisation by 3 dB
- [illegible]
- [illegible]

- $\gamma_f = 10$ dB

We then have a beacon power of ¼ Watt

This gives $P_{t\text{-}MAX} = 2.5 \times 10^{-5}$ W/Hz

Thus, in this example, a secondary transmitter with a bandwidth, for example, of 100 KHz would be allowed to transmit up to 2.5 W if it could not hear a beacon. As the beacon was detected and its received power climbed above L_S then the maximum allowable power would be reduced pro-rata.

Several interesting points arise from this simple analysis

1. The appropriate beacon power is independent of
    a. Primary system bandwidth
    b. Operating range
    c. Range to the secondary transmitter(s)
    d. Primary system required signal to noise ratio
2. The beacon power depends only on the primary receiver noise figure and the allowable de-sensitisation. Once a reference for beacon power has been legislated (the above is only a suggestion), only relatively small variations will result. The reference power (assuming receiver noise temperatures no lower than 290° and desensitisation by no less than 3 dB) will be the maximum beacon power ever used.
3. Feeder loss and antenna patterns do not affect the beacon power or maximum secondary power conditions for the beacon system. They will, of course, affect the spectrum sharing that is available.
4. The maximum power a secondary transmitter may transmit depends on the sensitivity of its beacon receiver. Thus the equipment manufacturer is motivated to produce beacon receivers with low noise figure and good beacon detectability.

2 Beacon Band Structure

A beacon signal represents the reception of a particular frequency in a band. We need at least one beacon type for every frequency. The first questions are...

1. How large should a beacon managed band be?
2. How much spectrum should be allocated to beacon transmission?
3. How many frequencies should be allocated within a beacon managed band?

As in all tradeoffs the result will be a compromise. On one hand the beacon managed band should be a large as possible to reduce the relative overhead of the beacon channel. On the other hand it should be as narrow as possible in order to guarantee the same radio propagation characteristics across the band and, in particular, between any sub band and the beacon transmission band. At this point we recognise that is will not be practical to reduce the managed bandwidth down to the multipath fading correlation bandwidth of the channel. Thus the beacon transmission band must be wide enough to have reasonable inherent multipath diversity. If this beacon reception diversity is greater than that available to the users of the band then the interference sensing mechanism will be conservative – i.e. it will tend to overestimate the interference generated. In the alternative situation an underestimate may occur. It may be necessary to introduce a fading margin to account for this difference. This could be done by transmitting the beacon at slightly higher power.

Based on the above it is proposed that the beacon managed band should be no greater than about 5% of its centre frequency – this should give reasonable correlation of average propagation characteristics. Some care will need to be taken to ensure that no tight frequency selectivities in the antennas cause difficulties. Such a problem is more likely to arise in the primary service which will probably be incumbent. A secondary service should not have the same problem since it will probably be designed to use any of the frequencies in the beacon managed band. This problem could be overcome by biasing the beacon transmitted power according to the relative antenna gains in the beacon transmission band and the represented signal frequency.

Equally, the beacon transmission band should probably not exceed about 5% of the beacon managed band. This avoids unacceptable overhead for the beacons.

Consideration must be given to the provision of duplexing filters since the primary receiver must transmit and the secondary transmitter must receive. Clearly if the beacon must be transmitted for a frequency that is right next to the beacon transmission band there will be a problem. This difficulty can be solved by having two beacon transmission bands, one at each end of the beacon managed band. Frequencies usage is then represented by the further beacon transmission band. At first this might seem to double the overhead. However, this can be overcome by having the beacon transmission bands manage bands on either side. This is illustrated in Figure 2.

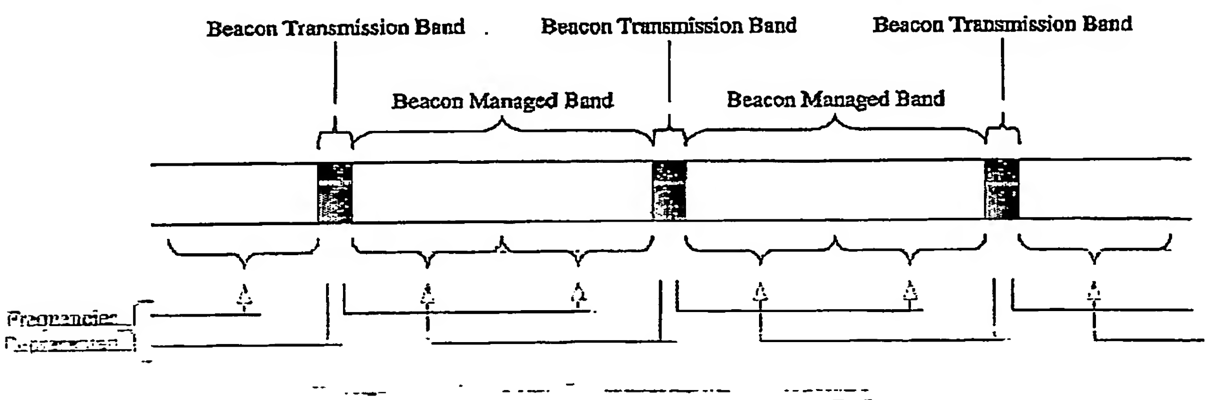

As an example, consider a band around 2 GHz. This could have a bandwidth of 100 MHz. This would lead to a beacon transmission band of 5 MHz (based on 5% bandwidth). According to the structure of Figure 2 there would always be a frequency spacing of at least 50 MHz to allow for duplexer filtering.

The above go some way towards answering the first two questions. The third question about the number of frequencies represented is again, a trade-off between complexity and flexibility. From the viewpoint of simplicity we would choose to have as few as possible. For example, if our 100 MHz wide band had an incumbent system with 5 MHz channels then the obvious answer would appear to be to have 20-off 5 MHz channels. However, this pre-supposes that there will be only one further tier of spectrum sharing and that the sharing system will also have a bandwidth of 5 MHz or multiples thereof. It may be that some sharing systems are introduced based on 7 MHz. In this case they may need to transmit as many as 3 beacons to keep a channel free. This would reserve 15 MHz of spectrum which would be inefficient. From the viewpoint of flexibility the answer would be to use the highest common denominator of all anticipated channel spacings. The disadvantage of this approach is that it might lead to the need for a large number of beacons for each used frequency. For example, a spacing of 1 MHz would require 5 beacons to keep a 5 MHz channel free. The number of frequencies would thus be a compromise between wasted spectrum and having a large number of beacons. At this stage a bandwidth of 5 MHz nevertheless seems reasonable leading to 20 beacons in our example.

3 Multiple Access Technology and Beacon Modulation Format

At any given beacon receiver we may have beacons from a number of receivers, relating to some different and some common frequencies. A would-be transmitter wants to examine beacons for a particular frequency. If there are many beacons for that frequency it is desirable that the task should not grow unduly. The options for multiple access are division in frequency, time and code (FDMA, TDMA and CDMA respectively). We consider these in turn

- FDMA – It is essential that the technology for beacon transmitters and receivers be low cost. This implies relatively inaccurate frequency references. For example 20 ppm end to end error would not be unreasonable. At 2 GHz this would correspond to an error of 40 kHz. This would make any FDMA system difficult to operate. Additionally, each carrier would be relatively narrowband, for example, representing 20 frequencies in 5 MHz would give a maximum channel spacing of 250 kHz. Thus we would not be exploiting the multipath diversity inherent in a bandwidth of 5 MHz.

- TDMA – There can practically be no co-ordination between the different beacon transmitters so no framing structures could be imposed. The use of GPS receivers to achieve this would be unattractive, particularly given the need for indoor operation. Thus the only available form of multiplexing in the time domain is random transmission akin to the Aloha protocol. One problem with this is that *a priori* we cannot be sure how many beacons we may need to receive in a given area so it will be difficult to dimension the average duty cycle of the transmission.

- CDMA – This would allow for a very large number of beacons relating to the same frequency since they would have the same code and would be separated only in time. One

problem with this, however, is the near far problem. If a strong beacon relating to one frequency is received then the receiver may know that that frequency is unavailable. However, it cannot know whether other frequencies are available since the beacon receiver is de-sensitised and it must be assumed that the other beacons *may* be present.

The overall conclusion from the above is that a combination of random TDMA and CDMA is required. The CDMA element will provide the ability to measure multiple beacons relating to the same frequency. The TDMA element will provide resistance to the near far effect. By listening for long enough, a beacon receiver can accumulate adequate confidence that, if there had been a beacon for a particular receiver at a lower level than a strong signal then it would have had an opportunity to receive it. Where a beacon transmitter needs to send multiple signals from one location (e.g. for a frequency that requires several contiguous sub bands) they must all be transmitted simultaneously and not in sequence. This is to avoid a single source occupying too much of the time.

The detailed structure of the code is beyond the scope of this discussion but some generalities can be considered. We want to make the beacon receiver as sensitive as possible. We can do this by making the correlation period for the CDMA component as long as practically possible. The problem here is that correlating for a long period one is subject to the effects of the frequency error between the beacon transmitter and the receiver. For example, at 2 GHz we might have an end-to-end error of 40 kHz. The longest period over which one can usually correlate is one half cycle of the error frequency, in our case 12.5 μs. This leads to an effective bandwidth of 80 kHz which is too wide for sensitive reception. This can be radically improved by the use of an FFT correlator [X]. By choosing the number of FFT bins we can make the correlation period any length. The maximum useful correlation period depends on the correlation time of the channel. This varies from about 200 μs to 1 ms depending on the mobility of the transmitter, receiver or both. We could make the underlying code have a duration of around 200 μs but allow up to, say, 4 repetitions which may be used in a static receiver to improve sensitivity. This would give an effective bandwidth from 5 kHz down to 1.25 kHz allowing good sensitivity. The FFT size would be 16, 32 or 64, all practical with today's technology for a 10 MHz sampling rate. There would be some scalloping loss from the FFT but this could be improved by randomising the frequency over ± 0.5 bins.

There is a further issue for the beacon information. On the one hand we only need to know the frequency represented by the beacon. However, we do need to distinguish beacons transmitted by our own system otherwise we will never transmit!

The approach for this could be as follows. At any given time there will only be a relatively small number of systems sharing any given band. Thus, only a small number of codes relating to system will be needed. Eight bits could be probably be enough. We arrange for every beacon transmitter to append its system type identifier to each beacon. This would be transmitted using the applicable spread spectrum code. This part would not be related to the represented frequency but would be married to it by its fine frequency (i.e. the FFT bin selected) and its timing. [illegible]

As stated earlier, for a given represented frequency the power to use for the beacon is the largest of those received. However, in multipath, there may be several paths all relating to the same beacon transmission. Inevitably there will be some ambiguity in determining this situation but this can be resolved in most cases by...

- Only taking other multipath components from the same FFT bin output as that which identified the strongest output – In the absence of significant Doppler the multipath components from any given beacon transmitter should all fall on the same fine frequency.
- Only take multipath components from a narrow window around the largest peak – This should constrain the components taken to within the maximum delay span of the propagation medium
- Take a maximum number of paths to avoid including peaks of noise

4 System Aspects

In the discussion so far we have considered a primary and a secondary system in which the receivers of the primary system have beacon transmitters and the transmitters of the secondary system have beacon receivers. In this case there is no protection for the secondary system. A more general set of possibilities is shown in Table 1.

Table 1 Types of System

Type of System	Beacon Transmitter	Beacon Receiver
1. Primary (Dominant)	√	
2. co-operating	√	√
3. Tertiary (Unprotected)		√

In addition, beacon transmissions could be made responsive for systems that could tolerate brief loss of communications. In this case a receiver would only transmit a beacon when it encountered unacceptable interference. The interfering sources would then detect the beacon and either stop transmitting or reduce their power.

Note it might in some cases be desirable for type 2 systems to receive and transmit essentially simultaneously on the beacon channel. This would apply if the dynamics of spectrum sharing were such that it was not appropriate for the terminal to stop transmitting the beacon when it was making its own system transmissions. The random TDMA transmission format for the beacon would allow this to happen.

The choice of level of protection for a system would be a matter for a combination of common sense and legislation.

There is potentially a great deal more flexibility in the actual beacon power a device may transmit. This may be influenced by a number of factors

- The manufacturer may need/wish to increase the nominal transmitted beacon power to take account of inaccuracies in RF gains, beacon transmitters etc. This would be an issue

for careful consideration for the regulator since one would not want to allow manufacturers to degrade their equipment performance requirements excessively at the expense of spectral efficiency. There could be a case for regulating an absolute maximum beacon transmitted power for a certain class of receiver.

- It may be desirable, and in some cases, acceptable to increase the beacon transmitted power to take account of multipath fading that may sometimes reduce the received beacon power without affecting the interference path

In addition, there are several fundamental approaches that could be applied to the setting of the beacon power...

1. Transmit fixed beacon power based on the equipment design – This approach is robust but not very flexible and may not lead to the best protection for the receiving equipment against multiple interferers

2. Arrange for the primary receiver to estimate the received noise and set the beacon power accordingly – In this case, as more interferers share the frequency the interference will rise. This rise will reduce the margin for further increase in interference leading to an increase in the beacon power. This increase should force the interferers to reduce their power (or prevent additional interferers from operating) in order to maintain the receiver sensitivity.

3. Arrange for the primary receiver to set its beacon power according to the maximum acceptable level of interference given the signal that it is receiving. This approach clearly should not be used in conjunction with power control within the system, as the two systems will interact in complex and unpredictable ways. However, for fixed power systems this approach could be used to maximise spectrum sharing.

The fundamental beacon concept is flexible enough to allow any or all of the above approaches to apply for different systems operating in the same beacon managed band.

In any given band, care would be needed to consider the types of system sharing the spectrum to ensure that overall stability is preserved. It would be possible, for example, to operate one of each of the types of system in Table 1 without instability. The type 1 system would not respond to beacons; the type 2 system would defer to the type 1 system and protect itself from the type 3 system. The type 3 system would defer to the types 1 and 2 systems and use whatever spectrum it could find. For complex sharing combinations would need more detailed evaluation.

For broadcast systems there are further degrees of freedom. For example, if one wished to use the beacon concept to allow sharing with the television band then every receiver would need to be equipped with a beacon transmitter. Note, however, that a beacon would only need to be transmitted corresponding to the channel that the viewer was receiving at the time.

5 Practical Issues

Because of the strong desirability of transmitting the beacon through the same antenna as that used for reception (in order to share the antenna's properties), this precludes the use of simple masthead pre-amplifiers. There are a couple of options

- Arrange for the masthead pre-amplifier to be the source of the beacon transmission. The reference power would be chosen in such a way that the absolute power requirement is always modest

- Build a diplexer into the masthead pre-amplifier with gain for the beacon transmitter as well.

Although these requirements add complexity they do not create insuperable problems. Additional revenue from the freed up spectrum should more than cover such costs.

6 Built In Self Test

The whole concept of spectrum sharing through beacons depends crucially on the beacons being received whenever present with power above the presumed reference sensitivity. An equipment with a failed beacon receiver may transmit at maximum power on any frequency and cause arbitrary interference. For this reason it is crucial that the status of the beacon receivers is known at all times. For this reason it is necessary to have a built in self test for the beacon receivers in which a known beacon is injected periodically into the receiver at limit range level in order to test its function. Failure to receive the test beacon will prevent the equipment from being allowed to transmit.